SEARCH STRATEGY, SAMPLING, AND COMPETITION LAW

Saul Levmore[1]

* * *

Introduction

Search costs matter and are reflected in many areas of law. For example, most disclosure requirements economize on search costs. A homeowner who must disclose the presence of termites saves a potential buyer, and perhaps many such buyers, from spending money to search, or inspect, the property. Similarly, requirements to reveal expected miles per gallon, or risks posed by a drug, economize on search costs. But these examples point to simple strategies and costs that can be minimized or entirely avoided with some legal intervention. Law can do better and take account of more subtle things once sophisticated search strategies are understood. This Essay introduces such search strategies and their implications for law.

High search costs make for more monopolies, but how bad are a monopolist's high prices? How does one search for a job? How many options should law want a corporate fiduciary to seek before agreeing to a merger? At present, these and many other questions tend to be resolved without thinking about the costs of searching for better options. Law needs to take search costs into account, plus the strategies that can be used to minimize these costs, far more than it does at present. This Essay will introduce important ideas about search costs by exploring several search strategies and a variety of their applications. This broad-brush approach leaves open, perhaps for future work, the task of digging down into particular industries. The goal is to convince the reader, and the legal (and law and economics) community more generally, that search costs deserve more attention.

One important search strategy is optimal stopping theory. It aims to find the best choice in a string of alternatives when an option that is passed by will no longer be available—so the searcher must "stop" and seize a winner. It may be useful to imagine that the decision-maker has difficulty comparing more than two alternatives at once. The idea is that if one cannot return to previously sampled options, like interviewees or goods for sale, then the correct strategy is to sample the field and then accept the first alternative that is superior

[1] William B. Graham Distinguished Service Professor of Law at The University of Chicago Law School

to the best option that was (but is by assumption no longer) available in the sample set. Admittedly, this requires the decision-maker to know the size of the field. It might be determined by nature, by reasonable deadlines for applicants, or by the cost of inspecting alternatives. A key insight concerns determining the optimal size of the sample set. An interviewer offered eight prospects should inspect three (1/e of the available group to be precise[2]) and then stop when observing the first subsequent interviewee who is better than the best in that sample of three. The decision-maker is thus limited to pairwise comparisons, a feature that may be convenient to decision-makers who have trouble keeping many things in mind at once. This strategy avoids cycling if there are three or more decision-makers because the choice is between just two options in any one step. Moreover, there is no returning to options passed by, just as rules against reconsideration avoid evidence of cycling and indeterminacy in deliberative groups governed by the usual rules of order.

In a set with eight options, ABCDEFGH, A is compared to B, the winner is then compared to C, and then the best of that sample group is kept in mind. If the next option, D, is inferior to the best of the sample, then D is dismissed, and we move on to E. If E beats the best of the sample, then E is chosen, and the search ends. The goal is to find the best in the A–H group of eight. Given the assumptions, this strategy is superior to any other in finding the best choice. The strategy fails when the best is in the ABC sample group, and it is also likely to fail when the best applicant is found late in the group. Thus, G might be the best, but it was dismissed once E was discovered to be better than any in the ABC group. Note, with some relief and some relaxation of the stated goal of finding the very best in the available group, that E is quite likely to be the second best if it is not the very best. It is apparent that the method does not take the costs of these one-on-one inspections, or comparisons, into account. These search costs are of the kind that drive the discussion in this Essay.

Another important piece of search theory, what we can call Diamond-price-adjustments, does consider search costs and has developed quite apart from optimal stopping theory. More significantly, neither has taken much root in the law and economics of antitrust, mergers and acquisitions law, criminal law, or employment law, though these subjects seem like fertile grounds for search theory.

---

[2] 1/e = 1 / 2.71828 . . , or roughly .37. In the practical world of sampling this is often referred to as the 37% rule.

Economists, beginning with Peter Diamond, have observed that consumer search costs raise prices, but much of the focus has been on sellers' incentives to obfuscate, raise search costs once a buyer is within reach, and then raise prices.[3] A seller can be expected to charge a bit more once it recognizes that a roving buyer will accept a price that is somewhat higher than one likely to be available elsewhere because the seller knows that the buyer will take its search costs into account. On the other hand, higher prices will discourage search—which is to say purchases and the costs undertaken to make these purchases—in the first place.

One aim of this Essay is to examine how these two theories—optimal stopping and Diamond-price-adjustments—affect prices. In turn, the insights are likely to shed light on legal interventions in the interest of consumer welfare. This Essay begins by sidelining optimal stopping in Part I. Optimal stopping theory is reintroduced in Part V after some discussion of Diamond-price-adjustments in Part II as well as search costs more generally, and particularly the disinclination of a first mover to be disadvantaged by a kind of free riding that this Essay will call "uncompensated sampling." For example, an interviewer who has no intention of hiring the first two candidates but uses them to sample the available pool and then follow the optimal stopping strategy is imposing costs on the candidates in the sample. Eventually, this Essay expands these ideas by considering two-sided search problems in Part V; a buyer might be searching for better prices, while the seller is also searching for a buyer that will produce the greatest profit. This two-sided inspection of search costs is new to law and economics. The examination of search and price adjustments offers serious implications for antitrust and other law. An incorporation of optimal stopping theory does yet more. The idea is to see competition through a lens that incorporates multiple approaches that are not normally thought of in one exercise and that, even alone, have rarely

---

[3] *See* Glenn Ellison & Alexander Wolitzky, *A Search Cost Model of Obfuscation*, 43 RAND J. ECON. 391, 417 (2012). For related analyses, see Peter Diamond, *A Model of Price Adjustment*, 3 J. ECON. THEORY 109, 156–68 (1971). *See also* Michael Baye, John Morgan & Patrick Scholten, *Information, Search, and Price Dispersion*, IND. UNIV. (Dec. 2006), https://perma.cc/LZ4A-K83P. Once search costs are "dispersed," things become more complicated, and high search costs can lead to *lower* prices because higher costs reduce demand. *See* Jose L. Moraga-Gonzalez, Zsolt Sandor & Matthijs R. Wildenbeest, *Do Higher Search Costs Make the Markets Less Competitive?*, IND. UNIV. (Mar. 2013), https://perma.cc/ZMD5-URUM.

found their way into antitrust law and economics, not to mention consumer protection law, corporate law, and other areas where "shopping" is common and often regulated.

Setting optimal stopping strategies aside allows us to concentrate on other insights from search theory. The analysis here begins with a buyer's search strategy, as if the seller is not also searching among potential bidders, or buyers. Put differently, most transactions involve two-sided search processes; the buyer searches for the best price (not to mention other features), and the seller evaluates buyers to set prices and decide when to accept offers. Still, it will be useful at the outset to view search costs from just one perspective. When optimal stopping is incorporated late in this Essay, along with a two-sided analysis, the number of balls in the air is greatly increased, and conclusions are hazy. This may be why the existing literature generally treats the problem as one-sided. In the case of understanding a firm's hiring process, for example—a typical application of optimal stopping—the focus has been on the employer's decision in hiring from a set of applicants who can be interviewed. The analysis here does the same, at least at first, ignoring the idea that the typical applicant is *also* choosing how many employers to consider and when to accept an offer rather than continue searching. Similarly, as we will see in the context of competition law, it is a leap (and a goal here) to think about buyers' choosing how many sellers to consider, but a further (usually ignored) leap would consider a seller's shopping strategy with multiple potential buyers—even if it is only to "interview" them as to the price they would pay. The reader can already see that this Essay will leave room for further work. A largely one-sided analysis is sufficiently novel and complex.

A word of warning and caution is in order. Most of the arguments developed here are limited to markets where the buyer faces search costs and where a single variable, price, is what matters. The arguments range across different markets in a manner that will annoy readers who are accustomed to low search costs or who look for things other than price. Some people like shopping; others have sufficient faith in other shoppers, and in perfect markets, that they take the first available option at the announced price. The markets used here for illustration are disparate, to be sure, but the idea is to show that search costs and search strategies play an important role in many markets. For example, a classic search case with variable prices is what some of us experience in exiting a highway and searching for gasoline. Prices tend to be higher at the station closest to the exit

ramp, especially so if these stations do not also have a local clientele. But if you think of drivers as looking for a station that accepts their gas-company credit card, or who stop at the first station because someone always needs the restroom in a hurry, then sellers will compete along dimensions other than price, and the discussion that follows will apply only to very few markets. More interesting, it is surprising that services like Waze do not (yet) encourage users to enter observed gas prices, burger prices, and restroom cleanliness the way they enter news about cars on the shoulders. In any event, many shoppers do compare prices at some cost. Despite our online world, it can be difficult to find actual—often post-bargain—prices without engaging in costly search. Sellers have some incentive to increase search costs for captive customers, and on other occasions they try to promise reduced search costs. A familiar example is that the actual price of a new car—including the net price after a trade-in—normally requires going to a dealership and improving on the sticker price. Sampling several dealerships can be time consuming. Similarly, outfitting a renovated kitchen with new appliances is a costly enterprise with significant search costs if the buyer is looking for a good price.

Even choosing a law school to attend can involve search costs. An applicant can discover rankings quite easily, but it requires effort to discover post-scholarship prices. A significant percentage of entering students will pay the sticker price, but many will bargain for better prices if they have attractive offers from competing schools. These are very different markets, and yet all involve costly searching by many buyers. To add to the fun, most of these markets reflect a degree of two-sided searching. The applicant searches for a better price from competing schools, and the schools search for applicants of various types. Similarly, a high school football star searches among universities to attend, while a given university in search of a talented football quarterback does not need to pay for *many* of these. It is a two-sided search process. In sum, this Essay is about the relationship between search costs and competition, and its message is that law should look at competition through the lens of search costs. Prices and even hints of collusion may be less important than is commonly thought.

I. Rewarding the First Mover

Optimal stopping theory casts light on the advantages and disadvantages of a first mover, such as a bidder or other party that

acts first and sometimes provides information to those who follow. The first applicant to be interviewed by an employer is likely to be disadvantaged because the employer will want to see other applicants. The first bidder for a corporation might be disadvantaged if other potential acquirers make use of the information provided by the price offered by the first mover. On the other hand, and more generally, most sellers want to attract bidders who would otherwise be discouraged by the prospect of wasting search costs.

A store that carries large kitchen appliances observes customers looking around, and making more than one visit, before they are ready to make a significant purchase. The store absorbs a cost by serving in the "sample." In many markets this is of limited concern because the store that is sampled is at least as likely to benefit from the sampling role played by competitors who are used by other shoppers. Similarly, a homebuyer might see a house that looked pretty good and might let it pass by without buying or bidding for it because the buyer wants to develop some knowledge about comparable houses that are available. At some point the buyer who, following the conventional assumption in the literature, can only keep a couple of houses in mind at once, stops and chooses a house that is superior to the one most preferred in the "sample" of those examined earlier. This may be because the buyer cannot keep more than a few houses in mind. It may also be because the potential buyer knows that by passing on an opportunity, the house under consideration might be sold to someone else before it can be reconsidered. This is less likely in the case of a washing machine, but even there a model might disappear from the market. It is easy to object to assumptions embedded in this description of comparison shopping. A sophisticated buyer might, for example, try to pay a potential seller to keep an opportunity open. The buyer might seek a right of first refusal to ensure that the home (or even the dishwasher) will not be sold quickly while the buyer is searching for the perfect home at a great price. This is easier for houses than it is for yet larger "goods," like target corporations, even though those searches might be mostly about price from the perspective of another corporation or venture capitalist.

If one knows they are in the sample used by a firm that follows optimal stopping theory, or a store knows that the customer is bargaining simply to use the price offered in order to get a better price elsewhere, then one is wasting time and energy. It would be nice to be paid for the information provided to the buyer. In most markets this does not lead to an undersupply of information because one does not

know if they are in the sample or if they are seeing customers who have acquired (free of charge, minus search costs) information from competitors. Of course, the buyer usually knows they are simply sampling, acquiring information, or engaging in a kind of free riding. One party's sampling is another's suffering from uncompensated sampling.

Even in the simplest situations, there is a cost to revisiting as well as to acquiring new information. Going to a new store or returning to one seen previously—a step ignored by the basic optimal stopping model, but in practice often available—requires time and energy. Many sellers do not want to promise or bargain about a price over the phone, or online, because they think the prospective buyer will keep calling around or use the offered price to get a better offer at another store. Sellers' reluctance to disclose price by phone or online is especially the case with car dealers, houses, and other big-ticket items, like household appliances. The seller knows, of course, that it is often worth the buyer's time to shop around, and the seller wants to avoid uncompensated sampling – a kind of free riding – by potential buyers. Even in this era of online information, it is difficult to bargain without absorbing search costs. Sellers know not to lower prices without forcing the buyer to pay in some way for information; they do not want to offer a better price that simply enables the buyer to get a better deal elsewhere.

One complication, or side topic here, is that search costs might allow price discrimination. If some people take the offered price and others are known to compare or bargain, a seller may be skilled at distinguishing between the two types. One group might be subsidizing the other or at least increasing the number of competitors. This is something to bear in mind when looking for empirical evidence of much that follows here. If consolidation causes prices to rise in hospitals, or among sellers of appliances, we need to know whether the observed prices are those stated for one group (those who do not compare prices) or the other.

Uncompensated sampling is found on both sides of transactions. When consumers search, and especially so at low cost to them, sellers also face costs and experience sampling; buyers can be using them, and even other sellers can benefit from the information they provide to buyers. It would be interesting to see a seller charge for a quote, but that is hard to structure because sellers could charge and then name a

high price that a buyer would never accept. The seller will have received a fee, but the buyer will obtain nothing of value. No buyer would pay for such a useless service. Perhaps a seller, like Seller #2, could demand a fee in return for giving a quote that is lower than one that the buyer claims to have received from Seller #1. This is not a strategy yet observed—even in law school admissions, where admissions officers regularly complain that prospective students, seeking to "buy" a spot in the class at a lower price, ask School #2 to match or beat an offer made by School #1. School #2 could charge for its improved offer both to earn money and to avoid uncompensated sampling by an applicant who intends to take this information from #2 to bargain with School #1 or even #3. Law journals might do the same. They do limit the amount of time an offer is left open, and they might charge submission fees, but neither approach is the same as gaining compensation for making an offer in the first place—and allowing the submitter to use this information to get another offer.

There are similar examples in the world of corporate acquisitions. Once an acquirer signals interest, there is an increased chance that another bidder will appear and benefit the target's shareholders. First movers have invested in studying potential targets, and they do not want their efforts to be used by other bidders who do not compensate the first mover. Any rules here can be understood as part of what we call competition law, as well as corporate law. Bidder #1 tries to get lockups or other devices that discourage other bids, or they seek ways of being paid for their first bids and their identification of a worthy target. They might have bought stock in the target and are then rewarded for their earlier search when #2 buys their stock at a higher price. (This comes closer to the idea for law school admissions officers.) Several well-known corporate law cases, such as *Smith v. Van Gorkom* (Del. 1985), involve a bidder's trying to finalize a takeover before another, often emulating, acquirer can enter the field. Once the second arrives, and an auction is in progress, current law all but forbids the managers of the target from interfering. The rule might benefit shareholders of the target in question, but it can discourage searches for viable targets in the future. Indeed, much of the analysis here, aimed at conventional competition law, can be applied to corporate law, where the regulatory apparatus is a combination of Delaware (or other state) law along with federal law that often requires equal payments to target shareholders. It goes without saying that the more Bidder #2 gains information, at no cost, from #1, the first mover, there will be a disinclination to go first. In the long run less searching, which is to say fewer first movers, will make everyone worse off.

These kinds of search costs, or bidding costs, present problems in many markets, but it is not clear that law can be of much help. For example, architects put in effort to secure contracts. Clients would like to see competing plans and prices, but architects are hesitant to put in effort if they only have a modest chance of success. The problem is partly solved by inviting architects to submit preliminary plans and then to invite just two or three to be finalists. Still, major architects do not want to enter these competitions unless winning comes with a large reward, in terms of money or reputation, to compensate for the effort put into unsuccessful proposals. Some architects charge for plans they deliver. Should law schools simply charge when an admitted student asks that the scholarship awarded be reconsidered and increased? Most clients will not simply pay for all architectural proposals, because (again) firms might put insufficient effort into plans just to earn the fee. One solution comes in the form of a prize (rather than something resembling the reward given by intellectual property law); the winning architect might extract a price premium.

The same is true for construction contracts and many large government contracts. The customer wants to see price competition, but if it pays challengers to enter the competition, it might get nothing of value in return. The problem is not simply overcome by the client's promising an add-on reward to the winner, because the second-best bidder will offer a better price in as much as the effort it put into its proposal is now a sunk cost. In the case of law school admissions, applicants put little effort into the marginal application, so the problem is different. Still, the suggestion offered earlier might have some promise with respect to large contracts with substantial application costs. A buyer (or even a target corporation) that receives a first offer might stipulate that it invites competition; it will pay for any bid that is better than the earlier bids. One objection to this idea, or explanation for why it has not materialized in the market, is that it is then unattractive to be the first bidder. As with corporate takeovers and many job offers, the one who invests first is at risk of providing uncompensated sampling; others can benefit from the information the first mover provides. And yet it is difficult to charge for getting things started because the paying party will fear that that the first mover will do just enough to earn the promised fee. In some ways, this is a familiar market failure. It is solved in the case of investments in art and some scientific advancements with a monopoly awarded by law in the form of copyrights and patents. But when the earlier investment

serves a single user, it is difficult to describe a comparable legal solution.

II. Sellers' Responses to Buyers' Search Costs

Turning now to, and building upon, an important element of search theory, the Diamond Paradox, imagine that you go to Store A where the price for the appliance you want is $500. From what you have seen in the past, $500 seems like a good price, but you would like to do better. It is costly to travel to Store B and then to C to compare prices because it takes time and entails transportation costs. And, again, you're only shopping for price; the analysis below will allow for qualities other than price. In other words, it's not like most experiences when choosing a restaurant or house, nor like a corporate acquisition, and certainly not like dating, where it is rare for two options to be identical, and one must think about multiple attributes rather than (for instance) price alone. Store A knows about your shopping costs because the costs are similar across buyers, especially when going from seller to seller. Moreover, A is probably an expert when it comes to assessing the likelihood that you enjoy shopping and are price sensitive.

Store A has its own costs and would be willing to sell you what you want for $500, but it knows that it will cost you at least $20 to go to the next store, B, so A offers to sell at $520. You can save money, including time, by not searching further. You have some information, and you suspect that something close to $500 is the best price you can get, so you might buy at $520 in Store A because it will be a waste to keep shopping. Admittedly this is a significant assumption, and quite unlike the implicit assumption in the optimal stopping world. As already suggested, perhaps the buyer has some prior about prices and thinks of $500 as a decent price. Let's assume that both buyer and seller know that $500 is in the range of what the buyer expects for the appliance in question. You, as the buyer, are a bit disappointed with A's price of $520, and you might hope to get $500 at the next store, but added together with your search cost of $20, you might as well buy at A for $520. The seller is rational for upping the price from $500 to $520, and you are rational in buying at $520.

But there are many clever sellers out there. The next one you could have searched, Store B, is probably thinking the same way and might also sell for $520, knowing that buyers who come to B do not

want to waste another $20 going to yet another store. This is easiest to see if we assume, more realistically than other assumptions in this Essay, that B does not know whether it is the first store you have visited. Each additional trip costs another $20, and of course no seller really knows whether you have already been to competing stores. Besides, returning to A is also costly. There is, apparently (and perhaps inefficiently), no cost or legal relief to the misinformed seller when you "fraudulently" tell a seller about your having been to other stores. It has been suggested that sellers (but not buyers) could be required to tell buyers about prices obtained from previous buyers, but that is off the table here.

If all the sellers will do this, and the price is going to be $520, then Store A will think again and set the price at $540. You will see $540 at A, and know that the best you can do is spend $20 to go to the next store where it is likely to be $520, so there is no point in searching beyond Store A. You would go to B if you thought B would sell at $500 which was your (correct) prior estimate for the price at which sellers would be willing to sell, but B has already figured things out and upped the price from $500 to $520. You too have figured this out. With this in mind, Store B will now also charge $540, knowing that if you are already there, you will not waste $20 to go to C to find a price of $520.

You can see where this piece of search theory (or Game Theory) is going. If all the stores are this rational, and they know now that all charge $540, they will then reason again and charge $560. But if B is also likely to sell at $560, A might as well charge $580, because A knows that you know it's not worth spending $20 to save $20. $580 in hand is worth $560 in the (costly to get to) bush.

There is a limit to this thinking. As the price goes up, and assuming the usual downward sloping demand curve, fewer people want to buy the item. At some point, it is not worth it for sellers to raise prices because there is more profit to be made by having more buyers. It turns out that this is the price a monopolist, who owns all the stores in town, would charge. If you find this implausible, perhaps because of its underlying assumptions, it is good enough to see that there is a limit to the game of price increases. Just as it is for a monopolist, there is a point at which sellers trade off the gain from higher prices with the loss of consumers. The key point here is that

prices will rise because of search costs and eventually level off. Let's imagine that the demand curve is such that the price settles at $600.

Even with a great deal of competition, stores might well offer the same price that the monopolist would charge. And because they settle at $600, many people who would have been happy to buy at $520 (or more) from sellers, who would themselves have been willing to sell at $500, now go without the good. This is remarkable, because it seems like it does not matter to buyers whether there is competition or a monopoly. With search costs in mind, $600 is the price.

There is room for legal intervention here. From a social point of view, lowering search costs—perhaps by requiring sellers of some goods to post enforceable prices online so that buyers don't need to engage in costly search—might be a better thing for law to do than to spend resources on lawyers who will dismantle monopolies. Lucky for us, we live in a time when we can usually find prices online, so that search costs for some goods (e.g., on Amazon) are close to zero, but let's continue to assume that sellers are not bound by previously posted prices and that online shopping will not suffice in some markets. There is some reason to think that sellers who do not honor previously posted prices by adjusting downwards will lose in the market, but as every shopper knows this is a weak claim. For some goods, as with cars and kitchen appliances, many buyers want to see the good they contemplate buying, and many buyers can bargain prices down. For other expensive goods, like Lasik surgery, this is simply not possible; the seller needs to see the prospective patient and the buyer might want to meet the doctor and thus see part of the "good" being offered. There are still many industries in which buyers face search costs. It is fair to say that when there are search costs, competitive sellers will take advantage of these costs by ratcheting prices up because buyers will save by curtailing their searches.

III. Monopolist Pricing

And what if there is a monopolist, M, in town? All the stores are owned by M. If M's price is $600, there is obviously no point in traveling to the next store owned by M. This is familiar in locations, often near bridges and tunnels, with many gas stations. Where many stations are owned by a single firm, they often charge the same price, and this might be true whether there is some competition, or all are owned by one firm. When they are competitors, they might charge the

same price when they are located far apart because of search costs, but then they might relocate over time to be near one another and take advantage of the best locations, considering highway exit ramps and other traffic patterns. This is familiar in terms of Nash Equilibrium as well as Hotelling-inspired thinking. It is the monopolist that might have stations sensibly apart from a social point of view, but the price will still be that stubborn $600 (you can think in terms of cents for high-octane gas). Even M does not charge $601 because as the price rises, demand drops. M does not want to lose too many customers. A separate question is why the monopolist, or near-monopolist, often has so many gas stations near one another. The answer seems to be that this is a way of excluding competitors, and a problem for antitrust law.[4]

Before expanding the analysis to include more than one (the price) characteristic, it is useful to state the (now) obvious antitrust implication. If competition yields the monopolist's price ($600 here), then why bother to use law to discipline a monopolist or oligopolists, even if there is evidence of collusion as to price or the exclusion of new competitors? As a matter of antitrust law, there is no objection to sellers with matching prices, but it is occasion for inspection as to explicit collusion. The reduction in consumer surplus will be the same whether competitors reach the price of $600 or the monopolist does so. Indeed, the monopolist might be superior once we take the spacing of stores into account and consider travel and energy consumption. It is worth noting that there are cases where the monopolist introduces additional costs as it tries to influence the government to allow its monopoly position. The analysis here sets aside this rent-seeking piece and is limited to the more familiar objection to monopolies.

The question is whether there are many industries that fit the description offered here. The argument is that anticompetitive appliance stores or car dealerships may not be worth fighting because even true competitors in their stead would reach the same price and bring about the same reduction in consumer surplus. On the other hand, a monopoly threat with respect to auto manufacturing or grocery stores is probably different. In one case, the game is about product

[4] Much of the empirical literature on gas prices is about price dispersion. Various countries are studied and several show that one gains somewhat by traveling a mile or kilometer in the search for gas. I encourage readers to do some local investigation; it keeps children occupied on long car trips.

rather than price competition, and in the other, there are so many products that shoppers have other reasons to favor one store over others. The focus here is on some of the markets where buyers face substantial search costs. With this in mind, the discussion turns to cases with multiple product features.

IV. Beyond Price Alone

Imagine you're shopping for a box of chocolates to give as a gift. The chocolates come in various stores, but even in one store, perhaps in an airport's duty-free shop, you can find different prices, brands, calories, colors, packaging, varieties, and so forth. How do you choose? It is difficult though not exactly costly to consider all the attributes, and you might want to weigh their importance to you. The super-rational thing to do, at least if there is a single decision-maker to avoid cycling among three or more alternatives, is to make a spreadsheet and list the attributes that matter. You might give them appropriate weights, study all the packages, record the information, and assign scores. You then total the scores and choose the one with the highest score. We all know people who accept or reject job offers this way. They consider salaries, locations, opportunities for advancement, branding, vacation days offered, and so forth. It's not easy to assign weights and points, but even if this is possible for some people, the entire process takes a great deal of time, so in the end even this search runs into costs. Calculations of this kind are not unlike regressions in their search for key variables. In the end, a buyer (or a market) might learn that it prefers the box that came from Belgium with a mixture of dark and milk chocolates, so long as it is not $5 more than the alternatives. A buyer might do this and reduce the decision-making to a choice among two or three boxes and then choose the one with the lowest price. Inevitably, there will be omitted variables. Perhaps the recipient or the buyer likes dark chocolates only in the winter. The studious buyer (like a regression approach) did not take seasons into account.

In any event, as an individual standing in the shop, even if you value the wisdom of the crowd, you don't have access to data from past purchases and the pleasure it gave recipients. You might instead assign weights and points and tally things up on your own. But what will you do with all these numbers? One approach—though it comes with many variations in the literature on decision-making—is to make use of pairwise comparisons. Perhaps one choice defeats another when the two are compared. You can do all the possible paired comparisons

and choose the one that gets the highest score (the greatest number of victories). There might even be one that beats all the others in head-to-head competition. If there is no such dominant choice, it is often impossible to come up with a method that is guaranteed to satisfy several unobjectionable aims. If you are worried about interactions between the attributes, you can compare pairs with pairs and so forth. But let's just focus on the fact that even if we could agree on some scoring and tallying method, you will probably find your favorite method too time-consuming when shopping for chocolate. Remember that this is not the familiar optimal stopping problem, because the chocolates are on several nearby shelves, and you can return to one you previously observed. For the same reason, the chocolate shopping example is not a conventional search theory problem because it does not cost much to go to the next item. And it is not a problem where Arrow's Theorem promises no solid solution because you are just one decision-maker with stable preferences. Still, the overall cost of considering all the attributes of all the chocolates is very high, considering how little is at stake.

Whenever we face decisions that can be made super-rationally, but only at great cost of time and effort, it is likely that heuristics have been developed to enable practical and quick decision-making. This can be hard-wired or developed (quite rationally) and learned. The choice among chocolates offers a good opportunity to think about one such heuristic, or satisficing method, developed by psychologists and called "take-the-best." The problem in the airport, as described here, is that there are many variables. The suggestion is that you first think about these variables and decide which is most important to you. At step two, you "take-the-best" judged by this variable. The psychology literature likes the example of buying milk at a large grocery store. Economists are unlikely to be impressed unless the idea is that one variable is so much more important than all the others that utility is easily maximized, or at least satisfied, by simplifying the problem in this way.

But going along with the psychologists, there really are many choices, and one could compare fat content, brand names, price, expiration dates, cow breeds, lactose levels, and popularity (perhaps uncompensated sampling of the choices others have made). The suggestion is that one picks the most significant variable—the literature seems to like *price*, but of course each shopper is free to pick something else—and then takes the best (here it is the lowest price) when comparing just this variable. More realistically, perhaps families

have decided whether to get 1% or whole milk, and the shopper then chooses the lowest price within that category. Or maybe it's three factors. You look at half gallons, with 2% fat, and then choose the one with the lowest price. Again, the idea is not to compare all seven or ten variables that can be observed on the many cartons. Consumer behavior of this kind is important to competitors as well as local monopolists and can be thought of along with the analysis that follows.

In passing, we can intuit about monopoly behavior regarding options offered to buyers. In a competitive market, entrants appear to satisfy fringe tastes. A monopolist has less incentive to do so, unless fringe buyers will vanish. This is easily tested. If a vacation spot has three ice cream stores, they are predicted to offer more flavors than a monopolist in a comparable town some distance away. This is because we imagine someone who loves passion fruit to settle for strawberry when the former is unavailable. The consumer loss from the monopoly is difficult to measure because it is about choices rather than prices. It is, more generally, about reduced innovation, but that is a familiar topic in antitrust. For example, in the modern tech world, monopolists may reduce innovation or bring it about, as start-ups innovate and yearn to be bought out by dominant firms looking to fill profitable niches.

We can see many examples of take-the-best in practice. An applicant is choosing which university to attend. Having been accepted to several, she visits them and has subtly examined various rankings, football team quality, fame of the math department, and so forth. But there remain many variables. Admissions officers will tell you that when it comes down to several choices, many admitted students will choose the college where they experienced the best party or the best student host on the visit. Those of us who think that faculty and classes matter might try to put ourselves in a student's shoes. Some version of take-the-best is not irrational when there are so many variables and limited ability to assign meaningful points (in the absence of a well-functioning market for each attribute). Are we sure that the student who follows the spreadsheet method really does better?

But how does one (even a take-the-best advocate) identify the most significant variable when choosing chocolate or milk or other things. Surveys suggest various shortcuts, often related to price. For example, when choosing among red wines at a restaurant, some seem

to buy the least expensive bottle; many choose the one with the median price; some choose the second most expensive, and some the second from the bottom, especially when the searcher is footing the bill. Take-the-best leaves open the question of how people decide which variable is most important to them. In practice, people choose different variables and then have different strategies with respect to each variable (like lowest or median price). Sellers are aware of this, to be sure, and the marketing literature is full of ideas about how to position and advertise goods based on surveys and observations about buyers' shortcuts. Strangely, this literature largely ignores the fact that the supermarket or other property owner also knows about the value of the best shelf space and will adjust prices accordingly. If so, it is not clear why the marketing department is terribly useful. But returning to the buyer's behavior, it is surely the case that people modify take-the-best and look at two rather than one variable among many. And they may *not* know how to weigh these two variables. They probably compare two things based on a couple of variables and then take the winner. They might combine an expanded version of take-the-best (two variables rather than one) with some pairwise competition.

There are many choices when it comes to milk, chocolate, and wine precisely because if people follow the take-the-best strategy, or an expanded version of it, the strategy produces variation. This is different from the simplified version of search theory offered earlier. All the stores ratcheted up to a price of $600 in our earlier examples, and now, with multiple characteristics and consumers responding with a take-the-best strategy, we find that because people identify different attributes as the one most important to them, there is variety. Uniformity came with a search for the best price, and now variety comes with the need to evaluate many attributes. The searching process is significant here but not for the same reason as was true for gas stations and appliance stores. If there is a lesson for competition law, it is probably about variety rather than price.

Indeed, few consumer markets are quite like gas stations. Store A charged $600 for an appliance, but a rational consumer who is about to accept that monopoly-like price might suddenly prefer the same model in silver rather than white. Store A does not stock the product in a color other than white and black; how many units and colors to stock is itself an interesting and two-sided question. There is suddenly a reason to go to Store B to check out the colors offered there. Store A might try to discourage the search for another color by offering a color it has in stock at $580 (or even $500), though the seller cannot

be sure that the customer is not bluffing. A clever buyer might feign interest in a color other than the ones in stock. In sum, search theory and take-the-best work together and might explain some of the price variation we observe.

I have already expressed the view that competition law is probably overvalued in markets where search is likely to drift to the monopoly price. It is now also clear that variety can be more important than price. The monopolist has some incentive to offer variety because it fears that competitors will enter the market to satisfy preferences not served by the monopolist. This is a familiar argument about variety and innovation in competitive verses monopolistic markets. In any event, antitrust law should probably pay little attention to the markets for consumer products where search is costly and where price is the important variable.

V. Putting Optimal Stopping in the Mix

It's time to incorporate optimal stopping. Imagine that five gas stations are available as you leave a highway. They are in the order of ABCDE. You might choose A because it is nearest and offers the lowest search costs. But right now your time is not terribly valuable, so search costs are low. You figure that A might have high prices because A knows that the location nearest the exit ramp is most attractive to many buyers. Indeed, it is observable that the cost of commercial property decreases with distance from the exit ramp. Perhaps A charges $4.70 per gallon, but you are a tourist and do not know whether much better prices are available in this state. The stations are a quarter mile apart, so you need to drive on to see the prices at the other stations.[5] When the cost is low, optimal stopping theory seems like a better tool than what we have examined to this point. You proceed to visit B and take the AB sample (because two out of five is the right size sample, and the exit sign indicated five stations), and decide that you will now shop at the first station that beats the better price offered by A and B. You cannot return to A or B because you are on the wrong side of a divided road, and it would be very expensive to return.

[5] Online prices (on Waze, for example) are notoriously inaccurate where prices often fluctuate day to day.

If all the station owners know that there are many optimal stoppers, C might simply charge a bit less than the lowest price offered by A or B. C's owner can drive by A and B each day and then adjust the price charged at C. Perhaps A charges $4.70 and B charges $4.50. We might know that D and E could be charging $4.20 and $4.00, respectively, but now we are just thinking about an optimal-stopper-type driver. A and B remain in business and appeal to buyers with high search costs. If there are many smart shoppers, it is plausible that the best price for C to offer, halfway down the road, is $4.40. An optimal stopper will have sampled A and B, and will now buy at C. This intuition is enhanced if we think that many reasonably smart optimal stoppers do not want to be left with E alone. In fact, going all the way to E is better than stopping at C or even D, if the gas stations are like people randomly located in the waiting room for interviews. But gas stations are not like that—if E is last, and there are many determined shoppers, it will pay for E to have the highest price because searching (returning to B or D) is very costly. E might charge the most but get the least business. Setting aside the question of how D should then price, or whether E should expect a new competitor, F, to arrive yet further down the road, and so forth, there is a good case to be made for C to price in the manner already suggested.

What if all the stations are owned by a monopolist, M2? If people learn that M2 owns all the stations, they will know that a rational M2 has no reason to lower the price at its CDE stations, at least if getting back on the highway adds serious search costs. Again, M2 might have purchased and staffed the four unnecessary stations in order to eliminate competition. A clever M2 might offer prices of $4.70, $4.50, $4.70, $4.70, and $4.70. Some drivers will search and feel clever when they find $4.50. Others will be optimal stoppers and proceed until they are stuck at the end at $4.70. Even with such a rational M2, the gain from ousting the monopolist, who might have worked with landowners to exclude competitors, is not great. It is not the difference between the competitors A and E ($4.70 and $4.00 a gallon) but rather between A and C ($4.70 and $4.40) or even between $4.50 and $4.40.

This sort of skepticism about anticompetitive markets is different, and certainly weaker, if we move away from gas stations near highway exit ramps, and look at other markets, where we cannot be sure that the buyer will start with A. Buyers may be able to start randomly at A or B, all the way to E, when they drive about town to look at appliances. The sellers do not know whether they will be in the sample group. If a seller is in the sample group of two, there is nothing

to do; the buyer will sample and then find returning to the sample too costly. If a seller is first after the sample group, the seller will want to be slightly better (lower price) than the lowest in the sampled group. But this is hard to assess, and therefore hard to play our iterative game. If ABCD are pricing at $4.10, $4.30, $4.50, and $4.70, for example, and we focus on E's strategy along with appropriate search costs, the chance that the sample group, randomly chosen by the buyer, was AB is not very high. There is no need to choose $4.00 in order to be the lowest price. Similarly, $4.40 is unlikely to be the right choice because that requires the sample group to contain CD exactly (with prices of $4.50 and $4.70). Imagine E chooses $4.20 in order to fit between $4.10 and $4.30. If E is in the sample group, E knows that its decision does not matter. Note that this is better for consumers than what we saw without randomization regarding the starting point for the search process. But, again, defeating the monopolist gets a price of $4.20, not $4.00. Indeed, a very clever monopolist might price at $4.40, rather than at $4.70, so as not to look like a dangerous monopoly, but that is a game theory (and political) strategy beyond the reach of this Essay.

Yet more interesting is to begin with the earlier search theory analysis but then combine it with some optimal stopping theory. Perhaps most sellers contemplate buyers' search costs and settle at a price of $4.70 (or $600 if we return to our first example of the market for appliances). But many buyers now guided by optimal stopping theory, and facing relatively low search costs, might create a sample and look for the first seller that can beat the best price in the sample. If the sample, AB, shows $4.70, then all C needs to do is price at $4.60. The possibilities are endless and depend on search costs and the willingness of sellers to experiment and take risks. But here too a two-sided approach complicates things, because buyers might also take risks and search more than first seems rational. In any event, prices in the competitive market can easily come close to the monopolist's price.

Anti-monopoly effort seems even less valuable if we add in the buyer's search costs and allow for a two-sided analysis. If buyers value their time, they might sample just A or just one randomly chosen firm (and price) from the ABCDE group. The right sample size is now less than 1/e (times the size of the group) because of search costs. Now the best strategy for a seller is to price higher than before, and yet closer to the monopoly price. Put differently, it is more likely for the seller to be able to best the price offered by the smaller sample group.

Readers might want to apply this to the most familiar example of optimal stopping theory. First, if the interviewer should sample 1/e applicants, but now we vary the analysis to include the cost of interviewing each additional applicant, then the optimal sample group is smaller, and depends on the cost of undertaking each interview. Each firm will have different search costs, but for many firms a small sample group will be right.

Finally, what about the same search theory and optimal stopping theory thinking as applied on the seller's side in the case of other consumer goods, including housing? In many markets, the seller can sample consumers to assess the prices they are willing to pay. In some markets, sellers can do this through auctions – and with inexperienced buyers, the seller may even benefit from winner's curse. But where it is difficult to assemble a group of competing buyers, or where buyers know that the seller has many identical products to sell, auctions can be unattractive to sellers. The seller can start with a high price and then observe offers made by potential buyers. It is no surprise that this resembles a Dutch Auction. The seller in a conventional shop creates a sample of buyers, before putting items on "sale." This is not unlike the way an employer, looking to hire one person, might investigate a sample group of job applicants, before taking the first applicant who is superior to the best observed in the sample. Anyone who has sold a house (other than in localities where auctions among bidding buyers are the norm) is familiar with this sort of updating as offers come in from potential buyers. Even a monopolist needs to sample buyers to estimate the demand curve. Again, the larger message is that the difference between monopoly and competitive prices shrinks for many different reasons once we look at markets with search costs in mind.

Conclusion

It is apparent that there is room for competition law, as well as other areas of regulation, to be integrated with optimal stopping and other theories about search strategies. For now, even one modest claim might be startling: Under plausible conditions, a monopolist or collusive arrangement is not more, or is just a bit more, destructive to social welfare than a group of competitors responding cleverly to rational consumers who face search costs. Even if these are modest, once we think about buyers' and sellers' responses to the behavior

imputed to others, seemingly competitive markets begin to set prices closer to the monopoly prices that are normally expected. Law's focus might be better directed at reducing search costs than on hunting down apparent cartels. Meanwhile, as discussed earlier in this Essay, there is probably a market failure in the provision of information to those who search. This Essay has offered a potential solution for markets (like law school admissions) where participants' search costs are low, but uncompensated sampling threatens the provision of information and competition in the first place. Where search costs are high, as in corporate acquisitions, we should focus more on the incentives for socially productive searches than on maximizing the benefits to a given set of a target firm's shareholders.

* * *

Saul Levmore is the William B. Graham Distinguished Service Professor of Law at the University of Chicago Law School.